Rank of mapping tori and companion matrices

Gilbert Levitt and Vassilis Metaftsis

Abstract

Given $\varphi \in GL(d,\mathbb{Z})$, it is decidable whether the mapping torus $G = \mathbb{Z}^d \rtimes_\varphi \mathbb{Z}$ has rank 2 or not (i.e. whether G may be generated by two elements); when it does, one may classify generating pairs up to Nielsen equivalence. If φ has infinite order, the rank of $\mathbb{Z}^d \rtimes_{\varphi^n} \mathbb{Z}$ is at least 3 for all n large enough; equivalently, φ^n is not conjugate to a companion matrix in $GL(d,\mathbb{Z})$ if n is large.

1 Introduction

The rank of a finitely generated group is the minimum cardinality of a generating set. There are very few families of groups for which one knows how to compute the rank (see [7] and references therein), and there exists no algorithm computing the rank of a word-hyperbolic group [2].

By Grushko's theorem, rank is additive under free product. It does not behave as nicely under direct product, even when one of the factors is $\mathbb{Z}$: the solvable Baumslag-Solitar group $BS(1,2) = \langle a, t \mid tat^{-1} = a^2 \rangle$ and the product $BS(1,2) \times \mathbb{Z}$ both have rank 2.

In this paper we consider semi-direct products $G = A \rtimes_\varphi \mathbb{Z}$ (also known as mapping tori), with the generator of the cyclic group $\mathbb{Z}$ acting on A by some automorphism $\varphi \in Aut(A)$. This was motivated by the remark that, when A is a free group F_d and φ has finite order in $Out(F_d)$, then G is a generalized Baumslag-Solitar group and its rank may be computed [10]. But we do not know how to compute the rank when φ has infinite order. Abelianizing does not help much, so we ask:

Question. *Given $\varphi \in GL(d,\mathbb{Z})$, can one compute the rank of $G = \mathbb{Z}^d \rtimes_\varphi \mathbb{Z}$?*

We can prove:

Theorem 1.1. *Given $\varphi \in GL(d,\mathbb{Z})$, one can decide whether $G = \mathbb{Z}^d \rtimes_\varphi \mathbb{Z}$ has rank 2 or not.*

It turns out that the rank of G is 1 plus the minimum number k such that $\mathbb{Z}^d$ may be generated by k orbits of φ (i.e. there exist $g_1, \dots, g_k \in \mathbb{Z}^d$ such that the elements $\varphi^n(g_i)$, for $n \in \mathbb{Z}$ and $i = 1, \dots, k$, generate $\mathbb{Z}^d$). In particular, G has rank 2 if and only if $\mathbb{Z}^d$ may be generated by a single φ-orbit. This happens precisely when φ is conjugate to the companion matrix having the same characteristic polynomial. This may be decided since the conjugacy problem is solvable in $GL(d, \mathbb{Z})$ [5].

Theorem 1.1 extends to the case when φ is an automorphism of an arbitrary finitely generated nilpotent group A.

When G has rank 2, one can classify generating pairs up to Nielsen equivalence. In particular:

Theorem 1.2. *Suppose that $G = \mathbb{Z}^d \rtimes_\varphi \mathbb{Z}$ has rank 2. There are infinitely many Nielsen classes of generating pairs if and only if the cyclic subgroup of $GL(d, \mathbb{Z})$ generated by φ has infinite index in its centralizer.*

Our next result is motivated by the following theorem due to J. Souto:

Theorem 1.3 ([11])**.** *Let A be the fundamental group of a closed orientable surface of genus $g \geq 2$. Let φ be an automorphism of A representing a pseudo-Anosov mapping class. Then there exists n_0 such that the rank of $G_n = A \rtimes_{\varphi^n} \mathbb{Z}$ is $2g+1$ for all $n \geq n_0$.*

We prove:

Theorem 1.4. *Given φ of infinite order in $GL(d, \mathbb{Z})$, with $d \geq 2$, there exists n_0 such that the rank of $G_n = \mathbb{Z}^d \rtimes_{\varphi^n} \mathbb{Z}$ is ≥ 3 for all $n \geq n_0$.*

The theorem becomes false if the hypothesis that φ has infinite order is dropped, or if 3 is replaced by 4. We do not know hypotheses that would guarantee that the rank is $d+1$ for n large.

An equivalent formulation of Theorem 1.4 is:

Theorem 1.5. *Given a matrix M of infinite order in $GL(d, \mathbb{Z})$, with $d \geq 2$, there exists n_0 such that M^n is not conjugate to a companion matrix if $n \geq n_0$.*

Our proof is based on the Skolem-Mahler-Lech theorem on linear recurrent sequences [3]. There are alternative approaches based on equations in S-units and Baker's theory on linear forms in logarithms. They are due to Amoroso-Zannier [1] and yield uniformity: *one may take $n_0 = [Cd^6(\log d)^6]$ where C is a universal constant (independent of M).*

We conclude with a few open questions.

Our analysis on $\mathbb{Z}^d$ uses the Cayley-Hamilton theorem. This is not available in a non-abelian free group F_d. Given $\varphi \in Aut(F_d)$, can one decide whether F_d may be generated by a single φ-orbit? More basically: given $\varphi \in Aut(F_d)$ and $g \in F_d$, can one decide whether the φ-orbit of g generates F_d?

What about ascending HNN extensions? For instance, let φ be an injective endomorphism of $\mathbb{Z}^d$ (a matrix with integral entries and non-zero determinant). Let $G = \mathbb{Z}^d *_\varphi = \langle \mathbb{Z}^d, t \mid tgt^{-1} = \varphi(g)\rangle$. Can one decide whether G has rank 2?

Acknowledgements. We wish to thank J.-L. Colliot-Thélène, F. Grunewald, P. de la Harpe, G. Henniart, and number theorists in Caen, in particular F. Amoroso, J.-P. Bezivin, D. Simon, for helpful conversations related to this work. The second author would also like to thank LMNO of Université de Caen for their hospitality during the preparation of the present work.

2 Generalities

Let A be a finitely generated group. The letters a, b, v will always denote elements of A. We denote by i_a the inner automorphism $v \mapsto ava^{-1}$.

Given $\varphi \in Aut(A)$, we let G be the mapping torus $G = A \rtimes_\varphi \mathbb{Z} = \langle A, t \mid tat^{-1} = \varphi(a)\rangle$. There is an exact sequence $1 \to A \to G \to \mathbb{Z} \to 1$. Up to isomorphism, G only depends on the image of φ in $Out(A)$. Any $g \in G$ has unique forms $at^n, t^n a'$ with $n \in \mathbb{Z}$.

If N is a characteristic subgroup of A, we denote by $\bar\varphi$ the automorphism induced on A/N. There is an exact sequence $1 \to N \to A \rtimes_\varphi \mathbb{Z} \to A/N \rtimes_{\bar\varphi} \mathbb{Z} \to 1$.

The rank $rk(G)$ is the minimum cardinality of a generating set. We let $vrk(G)$ be the minimum number of elements needed to generate a finite index subgroup: $vrk(G) = \inf_H rk(H)$ with the infimum taken over all subgroups of finite index.

Two generating sets are Nielsen equivalent if one can pass from one to the other by Nielsen operations: permuting the generators, replacing g_i by g_i^{-1} or $g_i g_j$. For instance, any generating set of $\mathbb{Z}$ is Nielsen equivalent to $\{0, \dots, 0, 1\}$ by the Euclidean algorithm.

The φ-orbit of $a \in A$ is $\{\varphi^n(a) \mid n \in \mathbb{Z}\}$. We denote by $OR(\varphi)$ the minimum number of φ-orbits needed to generate A. Clearly $OR(\varphi) \le rk(A)$. We also denote by $VOR(\varphi)$ the minimum number of φ-orbits needed to generate a finite index subgroup of A, so $VOR(\varphi) \le vrk(A)$.

Lemma 2.1. *Given $a, a_1, \dots, a_k \in A$, the intersection $A' = \langle a_1, \dots, a_k, at\rangle \cap A$ is generated by the $(i_a \circ \varphi)$-orbits of $a_1, \dots, a_k$.*

The $(i_a \circ \varphi)$-orbits of $a_1, \dots, a_k$ generate A if and only if $a_1, \dots, a_k, at$ generate G.

Proof. One has $(i_a \circ \varphi)^n(v) = (at)^n v (at)^{-n}$ for $v \in A$ and $n \in \mathbb{Z}$. This shows that the $(i_a \circ \varphi)$-orbit of a_i is contained in A'. Conversely, if $v \in A'$, write it in terms of $a_1, \dots, a_k, at$. The exponent sum of t is 0, so v is a product of elements of the form $(at)^n a_i (at)^{-n}$.

If $A' = A$, then $\langle a_1, \dots, a_k, at\rangle$ contains A and at, so equals G. $\square$

Corollary 2.2. $rk(G) = 1 + \min_{a\in A} OR(i_a \circ \varphi)$.

Proof. $\leq$ is clear. For the converse, use that any finite generating set of G is Nielsen equivalent to a set $\{a_1, \dots, a_k, at\}$ (Euclid's algorithm). $\square$

Corollary 2.3. $vrk(G) = 1 + \min_{a\in A, n\neq 0} VOR(i_a \circ \varphi^n)$.

Proof. If $n \neq 0$ and the $(i_a \circ \varphi^n)$-orbits of $a_1, \dots, a_k$ generate a finite index subgroup of A, the subgroup of G generated by $a_1, \dots, a_k, at^n$ has finite index because it maps onto $n\mathbb{Z}$ and it meets A in a subgroup of finite index.

Any finite subset of G generating a finite index subgroup is Nielsen equivalent to $\{a_1, \dots, a_k, at^n\}$ with $n \neq 0$, and the $(i_a \circ \varphi^n)$-orbits of $a_1, \dots, a_k$ generate a finite index subgroup of A. $\square$

Corollary 2.4. *Suppose that A is abelian.*

1. *$rk(G) = 1 + OR(\varphi)$ and $vrk(G) = 1 + VOR(\varphi)$.*
2. *G has rank ≤ 2 if and only if A is generated by a single φ-orbit. A pair (a_1, at) generates G if and only if the φ-orbit of a_1 generates A.*
3. *$vrk(G)$ is computable.*

Proof. i_a is the identity and $VOR(\varphi) \leq VOR(\varphi^n)$, so 1 follows from previous results. 2 is clear.

For 3, first suppose $A = \mathbb{Z}^d$. View φ as an automorphism of the vector space $\mathbb{Q}^d$. Then $VOR(\varphi)$ is the minimum number of φ-orbits needed to generate $\mathbb{Q}^d$. This is computable (it is the number of blocks in the rational canonical form of φ). If A has a torsion subgroup T, then $A/T \simeq \mathbb{Z}^d$ for some d. Let $\bar\varphi$ be the automorphism induced on $\mathbb{Z}^d$. Then $VOR(\varphi) = VOR(\bar\varphi)$ is computable. $\square$

3 Computability

Suppose $A = \mathbb{Z}^d$ with $d \geq 1$. We view $\varphi \in Aut(A)$ as an automorphism of $\mathbb{Z}^d$ or as a matrix in $GL(d, \mathbb{Z})$. Its companion matrix M_φ is the unique matrix of the form

$$\begin{pmatrix} 0 & & & & * \\ 1 & 0 & & & * \\ & \ddots & \ddots & & * \\ & & 1 & 0 & * \\ & & & 1 & * \end{pmatrix}$$

having the same characteristic polynomial as φ (the empty triangles are filled with 0's, and $*$ denotes an arbitrary integer).

Lemma 3.1. *Let $\varphi \in GL(d, \mathbb{Z})$, with $d \geq 1$.*

1. *The following are equivalent:*
   - *(a) $G = \mathbb{Z}^d \rtimes_\varphi \mathbb{Z}$ has rank 2;*
   - *(b) $\mathbb{Z}^d$ may be generated by a single φ-orbit;*
   - *(c) There exists $a \in \mathbb{Z}^d$ such that $\{a, \varphi(a), \dots, \varphi^{p-1}(a)\}$ is a basis of $\mathbb{Z}^d$.*
   - *(d) φ is conjugate to its companion matrix M_φ in $GL(d, \mathbb{Z})$.*
2. *Suppose that the φ-orbit of a generates $\mathbb{Z}^d$. Then the φ-orbit of b generates $\mathbb{Z}^d$ if and only if $b = h(a)$ where $h \in GL(d, \mathbb{Z})$ commutes with φ.*

Proof. We already know that (a) is equivalent to (b). If a is the first element of a basis of $\mathbb{Z}^d$ in which φ is represented by the matrix M_φ, then the basis is $\{a, \varphi(a), \dots, \varphi^{d-1}(a)\}$ and the φ-orbit of a generates $\mathbb{Z}^d$, so $(d) \Rightarrow (c) \Rightarrow (b)$.

Conversely, suppose that the φ-orbit of a generates $\mathbb{Z}^d$. By the Cayley-Hamilton theorem, $\mathbb{Z}^d$ is generated by $\{a, \varphi(a), \dots, \varphi^{d-1}(a)\}$. This set is a basis of $\mathbb{Z}^d$ in which φ is represented by M_φ. This proves 1.

To prove 2, suppose that h commutes with φ, and define $b = h(a)$. The image of the basis $\{a, \varphi(a), \dots, \varphi^{d-1}(a)\}$ by h is $\{b, \varphi(b), \dots, \varphi^{d-1}(b)\}$, so the orbit of b generates. Conversely, if the orbit of b generates, define h as the automorphism taking $\{a, \varphi(a), \dots, \varphi^{d-1}(a)\}$ to $\{b, \varphi(b), \dots, \varphi^{d-1}(b)\}$. It commutes with φ because M_φ represents φ in both bases. $\square$

Proposition 3.2. *If A is nilpotent, one can decide whether $G = A \rtimes_\varphi \mathbb{Z}$ has rank 2 or not.*

Proof. If $A = \mathbb{Z}^d$, one has to decide whether φ is conjugate to its companion matrix M_φ in $GL(d,\mathbb{Z})$. This is possible because the conjugacy problem is solvable in $GL(d,\mathbb{Z})$ by [5].

We now assume that A is abelian. It fits in an exact sequence $0 \to T \to A \to \mathbb{Z}^d \to 0$ with T finite. We denote by $a \mapsto \bar{a}$ the map $A \to \mathbb{Z}^d$, and by $h \mapsto \bar{h}$ the natural epimorphism $Aut(A) \to Aut(\mathbb{Z}^d)$. They each have finite kernel.

We have to decide whether A may be generated by a single φ-orbit. We first check whether the matrix of $\bar{\varphi}$ is conjugate to its companion matrix. If not, the answer to our question is no. If yes, [5] yields a conjugator and therefore an explicit $u \in \mathbb{Z}^d$ whose $\bar{\varphi}$-orbit generates $\mathbb{Z}^d$.

We claim that A may be generated by a single φ-orbit if and only if there exist $a \in A$ mapping onto u, and $\psi \in Aut(A)$ of the form $h\varphi h^{-1}$ with $h \in Aut(A)$ and $[\bar{h}, \bar{\varphi}] = 1$, such that the ψ-orbit of a generates A.

The "if" direction is clear. Conversely, suppose that the φ-orbit of b generates A. Then the $\bar{\varphi}$-orbit of $\bar{b}$ generates $\mathbb{Z}^d$, so by Lemma 3.1 there exists $\theta \in Aut(\mathbb{Z}^d)$ commuting with $\bar{\varphi}$ and mapping $\bar{b}$ to u. Let h be any lift of θ to $Aut(A)$. Defining $a = h(b)$ and $\psi = h\varphi h^{-1}$, it is easy to check that the ψ-orbit of a generates A. This proves the claim.

We now explain how to decide whether a and ψ as above exist. Note that a and ψ must belong to explicit finite sets: a belongs to the preimage A_u of u, and ψ belongs to the preimage X_φ of $\bar{\varphi}$ in $Aut(A)$.

By Theorem C of [5], the centralizer of $\bar{\varphi}$ in $Aut(\mathbb{Z}^d)$ is a finitely generated subgroup and one can compute a finite generating set. The same is true of $D = \{h \in Aut(A) \mid [\bar{h}, \bar{\varphi}] = 1\}$, so we can list the elements ψ in the orbit $D\varphi$ of φ for the action of D on X_φ by conjugation.

By the claim proved above, A may be generated by a single φ-orbit if and only if there exist $a \in A_u$ and $\psi \in D\varphi$ such that the ψ-orbit of a generates A. To decide this, we enumerate the pairs (a, ψ) with $a \in A_u$ and $\psi \in D\varphi$. For each pair, we consider the increasing sequence of subgroups $A_N = \langle \psi^{-N}(a), \dots, \psi^{-1}(a), a, \psi(a), \dots \psi^N(a) \rangle$. It stabilizes and we check whether $A_N = A$ for N large.

This completes the proof for A abelian. If A is nilpotent, let B be its abelianization and let $\rho : B \to B$ be the automorphism induced by φ. If $G_\varphi = A \rtimes_\varphi \mathbb{Z}$ has rank 2, so does its quotient $G_\rho = B \rtimes_\rho \mathbb{Z}$. Conversely, if G_ρ has rank 2, it is generated by t and some $b \in B$ whose ρ-orbit generates B. Let a be any lift of b to A. The subgroup of A generated by the φ-orbit of a maps surjectively to B, so equals A by a classical fact about nilpotent groups (see e.g. Theorem 2.2.3(d) of [8]). Thus G_φ has rank 2. □

Corollary 3.3. *If $A = \mathbb{Z}^2$ or $A = F_2$, one can compute the rank of G.*

Proof. The rank is 2 or 3, so this is clear from the proposition if $A = \mathbb{Z}^2$.

Recall that the natural map $Out(F_2) \to Out(\mathbb{Z}^2) = Aut(\mathbb{Z}^2)$ is an isomorphism (both groups are isomorphic to $GL(2, \mathbb{Z})$). Given $G = F_2 \rtimes_\varphi \mathbb{Z}$, let ρ be the image of φ in $Aut(\mathbb{Z}^2)$. Consider $G_\rho = \mathbb{Z}^2 \rtimes_\rho \mathbb{Z}$. We prove that G and G_ρ have the same rank.

Clearly $2 \leq rk(G_\rho) \leq rk(G) \leq 3$. If G_ρ has rank 2, Lemma 3.1 lets us assume that ρ is of the form $\begin{pmatrix} 0 & \pm 1 \\ 1 & n \end{pmatrix}$. Since G only depends on the class of φ in $Out(F_2)$, it is isomorphic to

$$\langle a, b, t \mid tat^{-1} = b, tbt^{-1} = a^{\pm 1} b^n \rangle,$$

so has rank 2. ☐

4 Nielsen equivalence

Proposition 4.1. *Suppose that A is abelian and $G = A \rtimes_\varphi \mathbb{Z}$ has rank 2.*

1. *Any generating pair of G is Nielsen equivalent to a pair (a, t) with $a \in A$.*

2. *Two generating pairs (a, t) and (b, t), with $a, b \in A$, are Nielsen equivalent if and only if b belongs to the φ-orbit of a or a^{-1}.*

Proof. Given $x, y \in A$, and n, write

$$(x, ty) \sim ((ty)^n x (ty)^{-n}, ty) = (\varphi^n(x), ty)$$

and

$$(x, ty) \sim (\varphi^n(x), ty) \sim (\varphi^n(x), ty\varphi^n(x)) \sim (x, ty\varphi^n(x)).$$

Every generating pair is equivalent to some (a, ty), with the φ-orbit of a generating A. But $(a, ty) \sim (a, ty\varphi^n(a))$ so by an easy induction $(a, ty) \sim (a, t)$. This proves 1.

If $b = \varphi^n(a^\varepsilon)$ with $\varepsilon = \pm 1$, then $(b, t) = (\varphi^n(a^\varepsilon), t) = (t^n a^\varepsilon t^{-n}, t) \sim (a, t)$. The converse follows from Theorem 2.1 of [6]. We give a proof for completeness. If $(b, t) \sim (a, t)$, we can write $b = w(a, t)$ with w a primitive word with exponent sum 0 in t. Such a word is conjugate to $a^{\pm 1}$ in the free group $F(a, t)$, so b is conjugate to $a^{\pm 1}$ in G. Since A is abelian, b belongs to the φ-orbit of $a^{\pm 1}$. ☐

Remark 4.2. More generally, if A is abelian, any generating set of G is Nielsen equivalent to a set of the form $\{a_1, \dots, a_k, t\}$.

Remark 4.3. The proposition does not extend to nilpotent groups. Let A be the Heisenberg group $\langle a, b, c \mid [a, b] = c, [a, c] = [b, c] = 1\rangle$. Let φ map a to ab and b to b. The generating pairs (a, t) and (ac^{-1}, t) are Nielsen equivalent (even conjugate) but ac^{-1} does not belong to the φ-orbit of $a^{\pm 1}$. Moreover, (a, tc) is a generating pair which is not Nielsen equivalent to a pair (x, t) with $x \in A$. Indeed, if it were, then t would be conjugate to tca^k for some $k \in \mathbb{Z}$ by [6]. Counting exponent sum in a yields $k = 0$. But t and tc are not conjugate.

Corollary 4.4. *Let $A = \mathbb{Z}^d$. If G has rank 2, the number of Nielsen classes of generating pairs is equal to the index of the group generated by φ and $-Id$ in the centralizer of φ in $GL(d, \mathbb{Z})$.*

Proof. By Proposition 4.1 we need only consider generating pairs of the form (a, t). Fix one. To any $b \in \mathbb{Z}^d$ such that (b, t) generates G we associate the automorphism ψ_b of $\mathbb{Z}^d$ taking the basis $\{a, \varphi(a), \dots, \varphi^{d-1}(a)\}$ to the basis $\{b, \varphi(b), \dots, \varphi^{d-1}(b)\}$. By Lemma 3.1, the image of this map $b \mapsto \psi_b$ is the centralizer of φ in $GL(d, \mathbb{Z})$. By Proposition 4.1, $(b, t) \sim (a, t)$ if and only if ψ_b is $\pm\varphi^n$ for some $n \in \mathbb{Z}$. $\square$

Example. The number of Nielsen classes of generating pairs is always finite if $d = 2$. If $\varphi = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}$, this number is infinite.

5 Powers

Fix $\varphi \in GL(d, \mathbb{Z})$. Say that $v \in \mathbb{Z}^d$ is *φ-cyclic* if its φ-orbit generates $\mathbb{Z}^d$, or equivalently if $\{v, \varphi(v), \dots, \varphi^{d-1}(v)\}$ is a basis of $\mathbb{Z}^d$. The existence of such a v is equivalent to φ being conjugate to its companion matrix, and also to G having rank 2. If v is φ^n-cyclic for some $n \geq 2$, it is φ-cyclic since its φ^n-orbit is contained in its φ-orbit.

If v is φ-cyclic, we denote by δ_n the index of the subgroup of $\mathbb{Z}^d$ generated by the φ^n-orbit of v. It does not depend on the choice of v since φ always has matrix M_φ in the basis $\{v, \varphi(v), \dots, \varphi^{d-1}(v)\}$. Also note that $\delta_1 = 1$. The group $G_n = \mathbb{Z}^d \rtimes_{\varphi^n} \mathbb{Z}$ has rank 2 (equivalently, φ^n is conjugate to its companion matrix) if and only if $\delta_n = 1$.

Theorem 5.1. *If $\varphi \in GL(2,\mathbb{Z})$ has infinite order, the rank of $G_n = \mathbb{Z}^2 \rtimes_{\varphi^n} \mathbb{Z}$ is 3 for all $n \geq 3$.*

Proof. If G_n has rank 2 for some n, there exists a φ^n-cyclic element v. Such a v is also φ-cyclic. In the basis $\{v, \varphi(v)\}$, the matrix of φ has the form $M = \begin{pmatrix} 0 & \varepsilon \\ 1 & \tau \end{pmatrix}$ with $\varepsilon = \pm 1$. If finite, the index δ_n is the absolute value of the determinant c_n of the matrix expressing the family $\{v, \varphi^n(v)\}$ in the basis $\{v, \varphi(v)\}$. We prove the theorem by showing $|c_n| > 1$ for $n \geq 3$.

The number c_n is determined by the equation $M^n = c_n M + d_n I$. It follows from the Cayley-Hamilton theorem that the sequence c_n satisfies the recurrence relation $c_{n+2} - \tau c_{n+1} - \varepsilon c_n = 0$.

If $\varepsilon = -1$ one has

$$c_n = \prod_{k=1}^{n-1} (\tau - 2\cos\frac{k\pi}{n})$$

because c_n is a monic polynomial of degree $n-1$ in τ which vanishes for $\tau = 2\cos\frac{k\pi}{n}$ (one also has $c_n = U_{n-1}(\tau/2)$, with U_{n-1} a Chebyshev polynomial of the second kind).

If $\varepsilon = 1$ one has

$$c_n = \prod_{k=1}^{n-1} (\tau - 2i\cos\frac{k\pi}{n}).$$

Since φ is assumed to have infinite order, one has $\tau \neq 0$ if $\varepsilon = 1$, and $|\tau| \geq 2$ if $\varepsilon = -1$. One checks that $|c_n| > 1$ for $n \geq 3$ (for $n \geq 2$ if $\varepsilon = -1$). $\square$

Theorem 5.2. *Suppose that $\varphi \in GL(d,\mathbb{Z})$ has infinite order.*

1. *There exists n_0 such that $G_n = \mathbb{Z}^d \rtimes_{\varphi^n} \mathbb{Z}$ has rank ≥ 3 for every $n \geq n_0$. Equivalently: φ^n is not conjugate to its companion matrix for $n \geq n_0$.*

2. *More precisely, the minimum index of 2-generated subgroups of G_n goes to infinity with n.*

Note that there are arbitrarily large values of n for which the rank of G_n is $d+1$ (whenever φ^n is the identity modulo some prime number). As already mentioned, it is proved in [1] that n_0 may be chosen to depend only on d.

The key step in the proof of Theorem 5.2 is the following result.

Proposition 5.3. *If φ has infinite order and v is φ-cyclic, then the index δ_n of the subgroup of $\mathbb{Z}^d$ generated by the φ^n-orbit of v goes to infinity with n.*

Proof of the theorem from the proposition. As above, if G_n has rank 2 for some n, there exists a φ-cyclic element v. For n large one has $\delta_n > 1$, so G_n has rank > 2. Assertion 1 is proved.

For Assertion 2, suppose that there are arbitrarily large values of n such that G_n contains a 2-generated subgroup H_n of index $\leq C$, for some fixed C. This subgroup has a generating pair of the form (a_n, t_n) with $a_n \in \mathbb{Z}^d$, and the intersection of H_n with $\mathbb{Z}^d$ is generated by the φ^{nm_n}-orbit of a_n for some $m_n \geq 1$. It has index $\leq C$ in $\mathbb{Z}^d$.

The subgroup of $\mathbb{Z}^d$ generated by the φ-orbit of a_n has index $\leq C$, so we can assume that it does not depend on n. Call it J. It is φ-invariant so we can apply the proposition to the action of φ on J, with $v = a_n$. This gives the required contradiction. $\square$

Proof of Proposition 5.3. When $d = 2$, one easily checks that c_n, as computed above, goes to infinity with n. The proof in the general case is more involved.

Define numbers $u_k(i)$, for $k = 0, \dots, d-1$ and $i \geq 0$, by $\varphi^i(v) = \sum_{k=0}^{d-1} u_k(i)\varphi^k(v)$. The sequences $u_0, \dots, u_{d-1}$ form a basis for the space $\mathcal{S}$ of sequences satisfying the linear recurrence associated to the characteristic polynomial of φ (the recurrence is $\sum_{j=0}^{d} a_j u_k(i+j) = 0$ if the characteristic polynomial is $\sum_{j=0}^{d} a_j X^j$).

The index δ_n is the absolute value of the determinant c_n of the matrix $(u_k(ni))_{0\leq i,k\leq d-1}$ (it is infinite if the determinant is 0). We have to prove that, given $c \neq 0$, the set of n's such that $c_n = c$ is finite. We assume it is not and we work towards a contradiction.

A sequence satisfies a linear recurrence if and only if it is a finite sum of polynomials times exponentials, so c_n also is a recurrent sequence. The Skolem-Mahler-Lech theorem [3] then implies that $c_n = c$ for all n in an arithmetic progression $\mathbb{N}_0 \subset \mathbb{N}$.

We shall now replace the basis u_k of $\mathcal{S}$ by another basis w_k depending on the eigenvalues of φ. We then assume that $D_n := \det(w_k(ni))_{0\leq i,k\leq d-1} = c' \neq 0$ for $n \in \mathbb{N}_0$.

We order the eigenvalues λ_k of φ so that $0 < |\lambda_1| \leq |\lambda_2| \leq \cdots \leq |\lambda_d|$. First suppose that the eigenvalues are all distinct. We then choose $w_k(i) =$

$(\lambda_{k+1})^i$. In this case D_n is a Vandermonde determinant, for instance

$$D_n = \begin{vmatrix} 1 & 1 & 1 \\ (\lambda_1)^n & (\lambda_2)^n & (\lambda_3)^n \\ (\lambda_1)^{2n} & (\lambda_2)^{2n} & (\lambda_3)^{2n} \end{vmatrix}$$

for $d = 3$, so $D_n = \prod_{1 \le k < m \le d} \big((\lambda_m)^n - (\lambda_k)^n\big)$.

If all moduli $|\lambda_k|$ are distinct, then $|D_n|$ goes to infinity with n because its diagonal term

$$(\lambda_2)^n(\lambda_3)^{2n} \dots (\lambda_d)^{(d-1)n} = \Big(\lambda_2(\lambda_3)^2 \dots (\lambda_d)^{(d-1)}\Big)^n$$

has modulus bigger than all others.

If the λ_k's are distinct but their moduli are not, expand D_n as a sum $\sum_j \varepsilon_j {\mu_j}^n$ (with $\varepsilon_j = \pm 1$). Now there may be several (possibly cancelling) terms for which $|\mu_j|$ takes its maximal value $K = |\lambda_2(\lambda_3)^2 \dots (\lambda_d)^{(d-1)}|$. Note that $K > 1$ because otherwise all λ_k's have modulus 1, hence are roots of unity by a classical result, and φ has finite order.

Since $D_n = c'$ for $n \in \mathbb{N}_0$ and $K > 1$, one has $\sum_{|\mu_j|=K} \varepsilon_j {\mu_j}^n = 0$ for $n \in \mathbb{N}_0$. Call this sum $D_{n,K}$. Recall that $D_n = \prod_{1 \le k < m \le d} \big((\lambda_m)^n - (\lambda_k)^n\big)$. To expand this product, one chooses one of $(\lambda_m)^n$ or $(\lambda_k)^n$ for each couple k, m. The corresponding term contributes to $D_{n,K}$ if and only if one always chooses a term of maximal modulus. In other words, $D_{n,K} = \prod_{1 \le k < m \le p} E_{k,m}$ with $E_{k,m} = (\lambda_m)^n - (\lambda_k)^n$ if $|\lambda_m| = |\lambda_k|$ and $E_{k,m} = (\lambda_m)^n$ if $|\lambda_m| > |\lambda_k|$. Since the λ_k's are non-zero, $D_{n,K} = 0$ implies $(\lambda_k)^n = (\lambda_m)^n$ for some k, m with $k \ne m$, so that $D_n = 0$, a contradiction.

This completes the proof when the eigenvalues of φ are distinct. In the remaining case, the basis w_k must have a different form: if λ is an eigenvalue of multiplicity r, we use the sequences $\lambda^i, i\lambda^i, \dots, i^{r-1}\lambda^i$. For instance,

$$D_n = \begin{vmatrix} 1 & 0 & 0 & 1 \\ (\lambda_1)^n & n(\lambda_1)^n & n^2(\lambda_1)^n & (\lambda_4)^n \\ (\lambda_1)^{2n} & 2n(\lambda_1)^{2n} & (2n)^2(\lambda_1)^{2n} & (\lambda_4)^{2n} \\ (\lambda_1)^{3n} & 3n(\lambda_1)^{3n} & (3n)^2(\lambda_1)^{3n} & (\lambda_4)^{3n} \end{vmatrix}$$

when $d = 4$ and $\lambda_1 = \lambda_2 = \lambda_3 \ne \lambda_4$.

Calling $\nu_1, \ldots, \nu_q$ the distinct eigenvalues of φ, there exist integers a, b, c_k, d_{mk} (depending only on the multiplicities of the eigenvalues) such that

$$D_n = an^b \prod_{k=1}^{q} (\nu_k)^{nc_k} \prod_{1 \leq k < m \leq q} \big((\nu_m)^n - (\nu_k)^n\big)^{d_{mk}}$$

(see [4] or Theorem 21 in [9]). For instance, D_n as displayed above equals $2n^3(\lambda_1)^{3n}((\lambda_4)^n - (\lambda_1)^n)^3$.

If $K > 1$, we conclude as in the previous case. If $K = 1$, all eigenvalues are roots of unity and $D_n = n^b E_n$ where E_n only takes finitely many values and $b > 0$ (an eigenvalue ν_j of multiplicity $r \geq 2$ contributes $1 + \cdots + (r-1)$ to b). Such a product cannot take a non-zero value infinitely often. $\square$

Corollary 5.4. *If A is abelian, and $\varphi \in Aut(A)$ has infinite order, then $G_n = A \rtimes_{\varphi^n} \mathbb{Z}$ has rank ≥ 3 for n large. The minimum index of 2-generated subgroups of G_n goes to infinity with n.*

This follows readily from Theorem 5.2, writing $A/T \sim \mathbb{Z}^d$ with T finite. The analogous result for nilpotent groups is false, as the following example shows. Let A be the Heisenberg group as in Remark 4.3. If φ maps a to bc, b to ac^2, and c to c^{-1}, then $\varphi^{2n+1}(a) = bc^{1-n}$, so G_{2n+1} has rank 2 since a and $\varphi^{2n+1}(a)$ generate A. The automorphism induced by φ on the abelianization of A has order 2.

References

[1] F.Amoroso, U. Zannier, in preparation.

[2] G. Baumslag, C.F. Miller III, H. Short, Unsolvable problems about small cancellation and word hyperbolic groups, Bull. London Math. Soc. 26 (1994), 97–101.

[3] G. Everest, A. van der Poorten, I. Shparlinski, T. Ward, Recurrence sequences, AMS Mathematical surveys and monographs 104, 2003.

[4] R.P. Flowe, G.A. Harris, A note on generalized Vandermonde determinants, SIAM J. Matrix Anal. Appl 14 (1993), 1146-1151.

[5] F. Grunewald, Solution of the conjugacy problem in certain arithmetic groups. Word problems, II (Conf. on Decision Problems in Algebra, Oxford, 1976), pp. 101–139, Stud. Logic Foundations Math., 95, North-Holland, Amsterdam-New York, 1980.

[6] M. Heusener, R. Weidmann, Generating pairs of 2-bridge knot groups, arXiv:0902.0799.

[7] I. Kapovich, R. Weidmann, Kleinian groups and the rank problem, Geometry and Topology 9 (2005), 375–402.

[8] E.I. Khukhro, Nilpotent groups and their automorphisms, de Gruyter expositions in mathematics 8, 1993.

[9] C. Krattenthaler Advanced determinant calculus. The Andrews Festschrift (Maratea, 1998). Sém. Lothar. Combin. 42 (1999), Art. B42q, 67 pp.

[10] G. Levitt, in preparation.

[11] J. Souto, The rank of the fundamental group of certain hyperbolic 3-manifolds fibering over the circle, in The Zieschang Gedenkschrift, Geometry and Topology Monographs, Vol. 14, 2008.

Gilbert Levitt
Laboratoire de Mathématiques Nicolas Oresme
Université de Caen et CNRS (UMR 6139)
BP 5186
F-14032 Caen Cedex
France
e-mail: levitt@math.unicaen.fr

Vassilis Metaftsis
University of the Aegean
Department of Mathematics
832 00 Karlovassi
Samos, Greece
e-mail: vmet@aegean.gr